Joint Inverse Covariances Estimation with Mutual Linear Structure

Ilya Soloveychik and Ami Wiesel,
Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel

arXiv:1511.06462v1 [stat.ML] 20 Nov 2015

***Abstract*—We consider the problem of joint estimation of structured inverse covariance matrices. We perform the estimation using groups of measurements with different covariances of the same unknown structure. Assuming the inverse covariances to span a low dimensional linear subspace in the space of symmetric matrices, our aim is to determine this structure. It is then utilized to improve the estimation of the inverse covariances. We propose a novel optimization algorithm discovering and exploiting the underlying structure and provide its efficient implementation. Numerical simulations are presented to illustrate the performance benefits of the proposed algorithm.**

***Index Terms*—Structured inverse covariance estimation, joint inverse covariance estimation, graphical models.**

I. Introduction

Large scale covariance and inverse covariance estimation using a small number of measurements is a fundamental problem in modern multivariate statistics. In many applications, e.g., linear array processing, climatology, spectroscopy, and longitudinal data analysis statistical properties of variables can be related due to natural physical features of the systems. Such relations often imply linear structure in the corresponding population covariance or inverse covariance matrices.

In this paper we focus on structured inverse covariance (concentration) matrix estimation. There are plenty of examples of inverse covariance estimation with linear structure. A partial list includes banded [1]–[3], circulant [4], [5], sparse (graphical) models [6], [7], etc. An important common feature of these works is that they consider a single and static environment where the structure of the true concentration matrix, or at least the class of structures, as in the sparse case, is known in advance. Often, this is not the case and techniques are needed to learn the structure from the observations. A typical approach is to consider multiple datasets sharing a similar structure but non homogeneous environments [8]–[12]. This is, for example, the case in covariance estimation for classification across multiple classes [13]. A related problem addresses tracking a time varying covariance throughout a stream of data [14], [15], where it is assumed that the structure changes at a slower rate than the covariances themselves [16]. Here too, it is natural to divide this stream of data into independent blocks of measurements.

This work was partially supported by the Intel Collaboration Research Institute for Computational Intelligence, the Kaete Klausner Scholarship and ISF Grant 786/11.

A part of results was presented at the 23-th European Signal Processing Conference (EUSIPCO), August 31- September 4, 2015, Nice, France.

Our goal is to first rigorously state the problem of joint concentration matrices estimation with linear structure and derive its fundamental lower performance bounds. Secondly, we propose and analyze a new algorithm of learning and exploring this structure to improve estimation of the inverse covariance matrices. More exactly, given a few groups of measurements having different covariance matrices each, our target is to determine the underlying low dimensional linear space containing or approximately containing the concentration matrices of all the groups. The discovered subspace can be further used to improve the inverse covariance estimation by projecting any unconstrained estimator on it. Most of the previous works considered particular cases of this method, e.g. factor models, entry-wise linear structures like in sparse and banded cases, or specific patterns like in banded, circulant and other models. We propose a new generic algorithm based on the idea of Lasso regression, [17]. Namely, our method seeks for minimizing the joint negative log-likelihood penalized by the dimensionality of the subspace containing the inverse covariances.

In [18] we considered a similar setting, where the covariance matrices, rather than their inverses, share the same linear structure. There we approached the problem of low dimensional covariance estimation using the Truncated SVD (TSVD) technique applied to the sample covariance matrices (SCM-s) of the groups of measurements. The reason we do not use the TSVD approach in the case of concentrations is that it would require inversion of the SCM-s, which becomes problematic when the number of samples in groups is relatively small. The algorithm we propose here avoids this restriction and demonstrates good performance even when the amount of measurements is scarce. Due to its Lasso-type construction, our algorithm naturally involves regularization parameter, which can be tuned using cross-validation schemes. Instead, we propose a simple closed form and data dependent tuning approach.

The rest of the text is organized as following: first we introduce notations, state the problem and illustrate examples. Then we derive the lower performance bound, propose our Joint Inverse Covariance Estimation (JICE) technique and an efficient numerical method computing it. In addition, we provide an upper performance bound for the JICE algorithm and demonstrate its application to the regularization parameter tuning. Finally, we perform numerical simulations demonstrating the advantages of the proposed algorithm.

Given $p \in \mathbb{N}$, denote by $\mathcal{S}(p)$ the $l = \frac{p(p+1)}{2}$ dimensional

linear space of $p \times p$ symmetric real matrices. $\mathbf{I}_d$ stands for the $d \times d$ identity matrix. For a square matrix $\mathbf{M}$, its Moore-Penrose generalized inverse is denoted by $\mathbf{M}^\dagger$. For any two matrices $\mathbf{M}$ and $\mathbf{P}$ we denote by $\mathbf{M} \otimes \mathbf{P}$ their tensor (Kronecker) product. $\sigma_1(\mathbf{R}) \geqslant \cdots \geqslant \sigma_r(\mathbf{R}) \geqslant 0$ stand for the singular values of a rectangular matrix $\mathbf{R}$ of rank not greater than r. If $\mathbf{M} \in \mathcal{S}(p)$, we denote its eigenvalues by $\lambda_1(\mathbf{M}) \geqslant \cdots \geqslant \lambda_p(\mathbf{M})$ and write $\mathbf{M} \succ 0$ if $\lambda_p(\mathbf{M}) > 0$. $\|\cdot\|_F$ denotes the Frobenius, $\|\cdot\|_2$ - the spectral and $\|\cdot\|_*$ - the nuclear (trace) norms of matrices, correspondingly. For any symmetric matrix $\mathbf{S}$, $\mathbf{s} = \operatorname{vech}(\mathbf{S})$ is a vector obtained by stacking the columns of the lower triangular part of $\mathbf{S}$ into a single column. In addition, given an l dimensional column vector $\mathbf{m}$ we denote by $\operatorname{mat}(\mathbf{m})$ the inverse operator constructing a $p \times p$ symmetric matrix such that $\operatorname{vech}(\operatorname{mat}(\mathbf{m})) = \mathbf{m}$. Due to this natural linear bijection below we often consider subsets of $\mathcal{S}(p)$ as subsets of $\mathbb{R}^l$. In addition, let $\operatorname{vec}(\mathbf{S})$ be a p^2 dimensional vector obtained by stacking the columns of $\mathbf{S}$, and denote by $\mathcal{I}$ its indices corresponding to the related elements of $\operatorname{vech}(\mathbf{S})$. Finally, we write $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ when $\mathbf{x}$ is a centered normal random vector with covariance $\mathbf{Q}$.

II. Problem Formulation and Examples

Consider a heterogeneous Gaussian model, namely, assume we are given $K \geqslant l = \frac{p(p+1)}{2}$ independent groups of real p dimensional normal random vectors

$$\mathbf{x}_k^i \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k), \quad i = 1, \ldots, n, \quad k = 1, \ldots, K, \tag{1}$$

with n i.i.d. (independent and identically distributed) samples in each group and covariances

$$\mathbf{Q}_k = \mathbb{E}[\mathbf{x}_k \mathbf{x}_k^T], \quad k = 1, \ldots, K. \tag{2}$$

We assume that the inverse covariances

$$\mathbf{T}_k = \mathbf{Q}_k^{-1} \succ 0, \quad k = 1, \ldots, K, \tag{3}$$

exist and span an r dimensional linear subspace of $\mathcal{S}(p)$. Our main goal is to estimate this subspace and use it to improve the concentration matrices estimation.

Let us list a few common linear subspaces which naturally appear in typical signal processing applications.

- **Diagonal**: The simplest example of a structured concentration matrix is a diagonal matrix. This is often the case when the noise vectors are uncorrelated or can be assumed such with great precision. In this case $r = p$. Note that the diagonal structure remains unaltered under inversion, thus making diagonal concentration equivalent to the diagonal covariance case, considered in [18].
- **Banded**: It is often reasonable to assume that the non-neighboring elements of a normal random vector are conditionally independent given all the other elements. Specifically, claiming that i-th element of the random vector is conditionally independent on the h-th if $|i-h| > b$ leads to the b-banded inverse covariance structure. The subspace of symmetric b-banded matrices constitutes an $r = \frac{(2p-b)(b+1)}{2}$ dimensional subspace inside $\mathcal{S}(p)$. Banded inverse covariance matrices are ubiquitous in graphical models, [6], [7].
- **Circulant**: The next common type of structured concentration matrices are symmetric circulant matrices, defined as

$$\mathbf{T} = \begin{pmatrix} t_1 & t_2 & t_3 & \ldots & t_p \\ t_p & t_1 & t_2 & \ldots & t_{p-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ t_2 & t_3 & t_4 & \ldots & t_1 \end{pmatrix}, \tag{4}$$

  with the natural symmetry conditions such as $t_p = t_2$, etc. Such matrices are typically used as approximations to Toeplitz matrices which are associated with signals obeying periodic stochastic properties, e.g. the yearly variation of temperature in a particular location. A special case of such processes are classical stationary processes, which are ubiquitous in engineering, [4], [5]. Symmetric circulant matrices constitute an $r = p/2$ dimensional subspace if p is even and $(p+1)/2$ if it is odd. Interestingly, like in the diagonal case, this structure does not change under inversion, [19], making the estimation problems in covariances and concentrations analogous, [18].
- **Sparse**: Sparse inverse covariance models generalize banded structures and are very common. In multivariate Gaussian distributions zero entries of the concentration reveal conditional independences. When graph representation is utilized to express the relations between the variates, zeros in inverse covariance are translated to missing edges in the graph making it sparse. Recently, Gaussian graphical models have attracted considerable attention due to developments in biology, medicine, neuroscience, compressed sensing and many other areas, [6], [7], [20], [21]. Unlike the previous examples, most of the literature treating graphical models does not assume the graph structure to be known in advance and aims at learning this structure.

In the following it will be convenient to use a single matrix notation for the multiple concentration matrices

$$\mathbf{t}_k = \operatorname{vech}(\mathbf{T}_k), \quad k = 1, \ldots, K, \tag{5}$$

$$\mathbf{Y} = [\mathbf{t}_1, \ldots, \mathbf{t}_K]. \tag{6}$$

Using this notation, the prior subspace knowledge discussed above is equivalent to a low-rank constraint

$$\mathbf{Y} = \mathbf{U}\mathbf{Z}, \tag{7}$$

where $\mathbf{U} \in \mathbb{R}^{l \times r}$ and $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_K] \in \mathbb{R}^{r \times K}$. Essentially our problem reduces to estimation of $\mathbf{Y}$ assuming it is low-rank. In the analysis we will assume r is known in advance, but the algorithm we propose recovers it from the data. Note that unlike [18], we assume a linear model in (7) and not affine, meaning that we do not separate the mean explicitly. The reason is that the additional intercept term complicated the derivations and have not led to significant gains in our numerical simulations.

III. Lower Performance Bounds

We begin by addressing the inherent lower performance bounds. For this purpose we use the Cramer-Rao Bound

(**CRB**) to lower bound the Mean Squared Error (**MSE**) of any unbiased estimator $\widehat{\mathbf{Y}}$ of $\mathbf{Y}$, defined as

$$\mathbf{MSE} = \mathbb{E}\left[\left\|\widehat{\mathbf{Y}} - \mathbf{Y}\right\|_F^2\right]. \tag{8}$$

The **MSE** is bounded from below by the trace of the corresponding **CRB** matrix. To compute this matrix, for each i we stack the measurements $\mathbf{x}_k^i$ into a single vector

$$\mathbf{x}^i = \begin{pmatrix} \mathbf{x}_1^i \\ \vdots \\ \mathbf{x}_K^i \end{pmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{T}^{-1}), \quad i = 1, \ldots, n. \tag{9}$$

where the block-diagonal matrix $\mathbf{T}$ is defined as

$$\begin{aligned} \mathbf{T}(\mathbf{U}, \mathbf{Z}) &= \operatorname{diag}\{\mathbf{T}_1, \ldots, \mathbf{T}_k\} \\ &= \operatorname{diag}\{\operatorname{mat}(\mathbf{U}\mathbf{z}_1), \ldots, \operatorname{mat}(\mathbf{U}\mathbf{z}_K)\}. \end{aligned} \tag{10}$$

Here the operator $\operatorname{diag}\{\mathbf{T}_1, \ldots, \mathbf{T}_K\}$ returns a block-diagonal matrix of size $pK \times pK$ with $\mathbf{T}_k$-s as its diagonal blocks. The Jacobian matrix of this parametrization reads as

$$\begin{aligned} \mathbf{J} = \frac{\partial \mathbf{T}}{\partial(\mathbf{U}, \mathbf{Z})} &= \begin{pmatrix} \frac{\partial \mathbf{t}_1}{\partial \mathbf{U}} & \frac{\partial \mathbf{t}_1}{\partial \mathbf{z}_1} & 0 & \ldots & 0 \\ \frac{\partial \mathbf{t}_2}{\partial \mathbf{U}} & 0 & \frac{\partial \mathbf{t}_2}{\partial \mathbf{z}_2} & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{t}_K}{\partial \mathbf{U}} & 0 & 0 & \ldots & \frac{\partial \mathbf{t}_K}{\partial \mathbf{z}_K} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{z}_1^T \otimes \mathbf{I}_l & \mathbf{U} & 0 & \ldots & 0 \\ \mathbf{z}_2^T \otimes \mathbf{I}_l & 0 & \mathbf{U} & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{z}_K^T \otimes \mathbf{I}_l & 0 & 0 & \ldots & \mathbf{U} \end{pmatrix} \in \mathbb{R}^{lK \times (lr + Kr)}, \end{aligned} \tag{11}$$

where we have used the following notation:

$$\frac{\partial \mathbf{t}_k}{\partial \mathbf{U}} = \left[\frac{\partial \mathbf{t}_k}{\partial \mathbf{u}_1} \; \frac{\partial \mathbf{t}_k}{\partial \mathbf{u}_2} \; \ldots \; \frac{\partial \mathbf{t}_k}{\partial \mathbf{u}_r}\right], \tag{12}$$

and the formulae

$$\frac{\partial \mathbf{t}_k}{\partial \mathbf{u}_j} = \frac{\partial \mathbf{U}\mathbf{z}_k}{\partial \mathbf{u}_j} = z_k^j \mathbf{I}_l, \quad \frac{\partial \mathbf{t}_k}{\partial \mathbf{z}_k} = \frac{\partial \mathbf{U}\mathbf{z}_k}{\partial \mathbf{z}_k} = \mathbf{U}. \tag{13}$$

Lemma 1. *(Lemma 1 from [18])*

$$\operatorname{rank}(\mathbf{J}) = lr + Kr - r^2 \tag{14}$$

This lemma implies that $\mathbf{J}$ is rank deficient, as

$$\operatorname{rank}(\mathbf{J}) = lr + Kr - r^2 \leqslant \min[lK, lr + Kr], \tag{15}$$

reflecting the fact that the parametrization of $\mathbf{T}$ or $\mathbf{Y}$ by the pair $(\mathbf{U}, \mathbf{Z})$ is unidentifiable. Indeed for any invertible matrix $\mathbf{A}$, the pair $(\mathbf{U}\mathbf{A}, \mathbf{A}^{-1}\mathbf{Z})$ fits as good. Due to this ambiguity the matrix $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ is singular and in order to compute the **CRB** we use the Moore-Penrose pseudo-inverse of $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ instead of inverse, as justified by [22]. Given n i.i.d. samples $\mathbf{x}^i, i = 1, \ldots, n$, we obtain

$$\mathbf{CRB} = \frac{1}{n}\mathbf{J}\,\mathbf{FIM}(\mathbf{U}, \mathbf{Z})^\dagger \mathbf{J}^T. \tag{16}$$

For the Gaussian population the matrix $\mathbf{FIM}(\mathbf{U}, \mathbf{Z})$ is given by

$$\mathbf{FIM}(\mathbf{U}, \mathbf{Z}) = \frac{1}{2}\mathbf{J}^T \operatorname{diag}\left\{\left[\mathbf{T}_k^{-1} \otimes \mathbf{T}_k^{-1}\right]_{\mathcal{I},\mathcal{I}}\right\}\mathbf{J}, \tag{17}$$

where $[\mathbf{M}]_{\mathcal{I},\mathcal{I}}$ is the square submatrix of $\mathbf{M}$ corresponding to the index set $\mathcal{I}$, defined in the notations section. The bound on the **MSE** is therefore given by

$$\begin{aligned} \mathbf{MSE} &\geqslant \operatorname{Tr}(\mathbf{CRB}) = \frac{1}{n}\operatorname{Tr}\left(\mathbf{FIM}(\mathbf{U}, \mathbf{Z})^\dagger \mathbf{J}^T \mathbf{J}\right) \\ &= \frac{2}{n}\operatorname{Tr}\left(\left[\mathbf{J}^T \operatorname{diag}\left\{\left[\mathbf{T}_k^{-1} \otimes \mathbf{T}_k^{-1}\right]_{\mathcal{I},\mathcal{I}}\right\}\mathbf{J}\right]^\dagger \mathbf{J}^T \mathbf{J}\right). \end{aligned} \tag{18}$$

To get more intuition on the dependence of the **MSE** on the model parameters, we bound it from below. Denote

$$\underline{\lambda} = \min_k \lambda_p(\mathbf{T}_k), \quad \overline{\lambda} = \max_k \lambda_1(\mathbf{T}_k), \tag{19}$$

to get the bound

$$\begin{aligned} \mathbf{MSE} &\geqslant \frac{2\underline{\lambda}^2}{n}\operatorname{Tr}\left(\left[\mathbf{J}^T\mathbf{J}\right]^\dagger \mathbf{J}^T\mathbf{J}\right) \\ &= \frac{2\underline{\lambda}^2}{n}\operatorname{rank}\mathbf{J} = \frac{2\underline{\lambda}^2}{n}(lr + Kr - r^2). \end{aligned} \tag{20}$$

As expected, the dependence on the model parameters is similar to that obtained in [18] for the joint structured covariance estimation and in [23] for the problem of low-rank matrix reconstruction.

IV. JICE Algorithm

We now introduce a novel algorithm based on convex penalized likelihood relaxation, whose purpose is to estimate the concentrations $\mathbf{T}_1, \ldots, \mathbf{T}_K$ utilizing the prior that they belong to a low dimensional subspace. We, therefore, refer to it as the Joint Inverse Covariance Estimation (JICE) algorithm. The most natural brute force approach would be to form the K SCM-s

$$\mathbf{S}_k = \frac{1}{n}\sum_{i=1}^n \mathbf{x}_i^k \mathbf{x}_i^{kT}, \quad k = 1, \ldots, K, \tag{21}$$

invert them and fit by a subspace of a small dimension. This can be done, for example, by the means of Principal Component Analysis (PCA). Such approach, applied to the SCM-s (rather than to their inverses), was proposed in [18] to treat the problem of joint covariance estimation with mutual linear structure. When the number of samples n in each group is smaller than or even close to the dimension p (the scenario we are mostly interested in), the inversion of the SCM-s would be impossible due to rank deficiency or would approximate the true inverse covariances poorly, thus causing the proposed PCA algorithm to fail. Instead, we suggest a different approach based on regularized likelihood maximization and its efficient implementation.

A. The Basic Algorithm

Our aim is to estimate the true inverse covariances, while simultaneously trying to keep the dimension of the space spanned by the estimators small. For this purpose we suggest

the following regularized average log-likelihood optimization program

$$[\widehat{\mathbf{T}}_1,\ldots,\widehat{\mathbf{T}}_K] = \underset{\widetilde{\mathbf{T}}_1,\ldots,\widetilde{\mathbf{T}}_K}{\operatorname{argmin}} \frac{1}{K}\sum_{k=1}^{K}\left[-\log|\widetilde{\mathbf{T}}_k| + \operatorname{Tr}\left(\mathbf{S}_k\widetilde{\mathbf{T}}_k\right)\right] + \eta\left\|\widetilde{\mathbf{Y}}\right\|_*, \tag{22}$$

where

$$\widetilde{\mathbf{Y}} = \left[\operatorname{vech}\left(\widetilde{\mathbf{T}}_1\right),\ldots,\operatorname{vech}\left(\widetilde{\mathbf{T}}_K\right)\right], \tag{23}$$

and η is a regularization parameter. Program (22) assumes the samples are Gaussian and aims at minimizing the average negative log-likelihood of the K groups of measurements and simultaneously enforces joint low dimensional structure on the concentration matrices. The latter is achieved by penalizing the nuclear norm of $\widetilde{\mathbf{Y}}$. Nuclear norm is a convex envelope of the counting measure on singular values of a matrix. Therefore, intuitively the purpose of the penalty term in (22) is to decrease the number of non-zero singular values, [24]. The optimization problem (22) can be viewed as the analog of the Lasso estimator [17], tailored to low-rank matrices as opposed to sparse vectors. Note that the program at hand is convex, and can be treated by many standard numerical solver, such as CVX, [25], [26].

Usually the value of the regularization parameter η is specified by the statistician based on his expert prior knowledge on the estimated quantities and the desired balance between bias and variance of the estimate. As a part of the theoretical results, in the next section we recommend a particular choice of this parameter guarantying closeness of the estimate to the unknown true inverse covariance matrix in Frobenius norm.

B. Extensions of the JICE

The target function (22) of the proposed JICE algorithm assumes the data to be normally distributed and is, therefore, based on the Gaussian negative log-likelihoods. However, a natural extension can be easily obtained by replacing these log-likelihoods by any other convex functions. More generally, the same relaxed penalized likelihood optimization technique for joint estimation of matrices with a common linear structure can be applied to any other estimator given as a solution to a convex optimization program. For example, a straight forward extension to a robust structured covariance estimation is achieved by adding a regularizing nuclear penalty term to a joint target constructed from multiple convexly relaxed M-estimators. An examples of such relaxation is the COCA estimator, [27].

As an additional option, when some convex prior structure $\mathcal{S}$ on the unknown inverse covariances is given, this information can be easily incorporated into the target by imposing a set of natural convex constraints $\widetilde{\mathbf{T}}_k \in \mathcal{S},\ k=1,\ldots,K$.

C. Bias Removal

The JICE algorithm suffers from bias since we replace the intractable rank counting penalty on the estimated matrix $\mathbf{Y}$ by the nuclear norm, which is a convex proxy of the counting measure on singular values. We suggest an additional step aimed at removing this bias. Denote the vectorized solution of (22) by

$$\widehat{\mathbf{Y}} = \left[\operatorname{vech}\left(\widehat{\mathbf{T}}_1\right),\ldots,\operatorname{vech}\left(\widehat{\mathbf{T}}_1\right)\right].$$

Analogously to the original parametrization (7), consider the decomposition

$$\widehat{\mathbf{Y}} = \widehat{\mathbf{U}}\widehat{\mathbf{Z}}, \tag{24}$$

where $\widehat{\mathbf{U}} \in \mathbb{R}^{l\times s}$ has orthonormal columns and $\widehat{\mathbf{Z}} \in \mathbb{R}^{s\times K}$. Such decomposition naturally suggests treating the columns of $\widehat{\mathbf{U}}$ as the basis vectors of the approximate low dimensional subspace of $\mathcal{S}(p)$. We get an improved concentration matrices estimator by minimizing the negative log-likelihoods over the subspace spanned by $\widehat{\mathbf{U}}$

$$\widehat{\mathbf{T}}'_k = \underset{\widetilde{\mathbf{T}}_k\subset\operatorname{span}(\widehat{\mathbf{U}})}{\operatorname{argmin}} -\log|\widetilde{\mathbf{T}}_k| + \operatorname{Tr}\left(\mathbf{S}_k\widetilde{\mathbf{T}}_k\right), \quad k=1,\ldots,K. \tag{25}$$

Remarkably, this additional bias removal step requires solving K decoupled programs that can be performed in parallel.

D. ADMM Implementation

In spite of the fact the JICE algorithm described above can be solved by any standard general purpose convex numerical library, one of our goals was to develop a stand alone implementation. Following [28], we propose an efficient solution to (22) based on Alternating Direction Method of Multipliers (ADMM). The idea behind the ADMM is to optimize the augmented target

$$f(\widetilde{\mathbf{T}}_1,\ldots,\widetilde{\mathbf{T}}_K,\widetilde{\mathbf{Y}}) = \frac{1}{K}\sum_{k=1}^{K}\left[-\log|\widetilde{\mathbf{T}}_k| + \operatorname{Tr}\left(\mathbf{S}_k\widetilde{\mathbf{T}}_k\right)\right] + \eta\left\|\widetilde{\mathbf{Y}}\right\|_* + \frac{\rho}{2}\frac{1}{K}\sum_{k=1}^{K}\left\|\operatorname{vech}\left(\widetilde{\mathbf{T}}_k\right) - \widetilde{\mathbf{Y}}_{:,k}\right\|^2, \tag{26}$$

with respect to its parameters $\widetilde{\mathbf{T}}_1,\ldots,\widetilde{\mathbf{T}}_K,\widetilde{\mathbf{Y}}$ under the constraint

$$\text{s.t.} \qquad \operatorname{vech}\left(\widetilde{\mathbf{T}}_k\right) = \widetilde{\mathbf{Y}}_{:,k},\ \ k=1,\ldots,K. \tag{27}$$

Here $\widetilde{\mathbf{Y}}_{:,k}$ denotes the k-th column of $\widetilde{\mathbf{Y}}$ and the constraints enforce consensus between the variables (note that when the constrains are satisfied, the target coincides with the original one, (22)). The solution is obtained via introduction of dual variables $\mathbf{U}_k \in \mathbb{R}^l$ and follows the iterative scheme (we replace iteration indices by arrows to simplify notations)

$$\widetilde{\mathbf{T}}_k \leftarrow \underset{\overline{\mathbf{T}}_k}{\operatorname{argmin}} -\log|\overline{\mathbf{T}}_k| + \operatorname{Tr}\left(\mathbf{S}_k\overline{\mathbf{T}}_k\right) + \frac{\rho}{2}\left\|\operatorname{vech}\left(\overline{\mathbf{T}}_k\right) - \widetilde{\mathbf{Y}}_{:,k} + \mathbf{U}_k\right\|^2,\ k=1,\ldots,K, \tag{28}$$

$$\widetilde{\mathbf{Y}} \leftarrow \underset{\overline{\mathbf{Y}}}{\operatorname{argmin}}\,\eta\left\|\overline{\mathbf{Y}}\right\|_* + \frac{\rho}{2K}\sum_{k=1}^{K}\left\|\operatorname{vech}\left(\widetilde{\mathbf{T}}_k\right) - \overline{\mathbf{Y}}_{:,k} + \mathbf{U}_k\right\|^2, \tag{29}$$

$$\mathbf{U}_k \leftarrow \mathbf{U}_k + \operatorname{vech}\left(\widetilde{\mathbf{T}}_k\right) - \widetilde{\mathbf{Y}}_{:,k},\ k=1,\ldots,K. \tag{30}$$

The main advantage of the proposed ADMM technique is that both updates in (28) and (29) have easily computable closed form solutions. Indeed, the first order optimality conditions for the $\widetilde{\mathbf{T}}_k$ update target (28) yield

$$\rho\overline{\mathbf{T}}_k - \overline{\mathbf{T}}_k^{-1} = \rho \operatorname{mat}\left(\widetilde{\mathbf{Y}}_{:,k} - \mathbf{U}_k\right) - \mathbf{S}_k. \tag{31}$$

Denote by $\mathbf{D}_k$ the orthogonal matrix diagonalizing the right-hand side of (31), then

$$\mathbf{D}_k^T\left(\rho \operatorname{mat}\left(\widetilde{\mathbf{Y}}_{:,k} - \mathbf{U}_k\right) - \mathbf{S}_k\right)\mathbf{D}_k = \mathbf{\Lambda}_k. \tag{32}$$

Multiply (31) by $\mathbf{D}_k^T$ on the left and by $\mathbf{D}_k$ on the right to get

$$\rho\underline{\mathbf{T}}_k - \underline{\mathbf{T}}_k^{-1} = \mathbf{\Lambda}_k, \tag{33}$$

where $\underline{\mathbf{T}}_k = \mathbf{D}_k^T\overline{\mathbf{T}}_k\mathbf{D}_k$. We can now construct a diagonal solution of this equation by solving the p quadratic equations involving the diagonal entries of $\underline{\mathbf{T}}_k$ of the form

$$\rho\underline{t}_k^j - \frac{1}{\underline{t}_k^j} = \lambda_k^j, \;\; j = 1, \ldots, p, \tag{34}$$

where $\underline{\mathbf{T}}_k = \operatorname{diag}\left\{\underline{t}_k^1, \ldots, \underline{t}_k^p\right\}$ and $\mathbf{\Lambda}_k = \operatorname{diag}\left\{\lambda_k^1, \ldots, \lambda_k^p\right\}$. Solving these quadratic equations and choosing positive roots, yields

$$\underline{t}_k = \frac{1}{2\rho}\left(\lambda_k + \sqrt{\lambda_k^2 + 4\rho}\right). \tag{35}$$

and finally,

$$\underline{\mathbf{T}}_k = \frac{1}{2\rho}\left(\mathbf{\Lambda}_k + \sqrt{\mathbf{\Lambda}_k^2 + 4\rho\mathbf{I}}\right). \tag{36}$$

The corresponding iteration becomes

$$\widetilde{\mathbf{T}}_k \leftarrow \frac{1}{2\rho}\mathbf{D}_k^T\left(\mathbf{\Lambda}_k + \sqrt{\mathbf{\Lambda}_k^2 + 4\rho\mathbf{I}}\right)\mathbf{D}_k, \tag{37}$$

which is guaranteed to be a positive definite matrix.

In order to derive a closed form solution for the second iterative step (29), we introduce the matrix singular value soft thresholding operator. Given a matrix $\mathbf{M}$ with the SVD

$$\mathbf{M} = \mathbf{U}\begin{pmatrix}\sigma_1 & 0 & \ldots \\ 0 & \sigma_2 & \ldots \\ \vdots & \vdots & \ddots\end{pmatrix}\mathbf{V}^T, \tag{38}$$

the singular values soft thresholding operator is defined as

$$S_\varepsilon(\mathbf{M}) = \mathbf{U}\begin{pmatrix}\max[\sigma_1 - \varepsilon, 0] & 0 & \ldots \\ 0 & \max[\sigma_2 - \varepsilon, 0] & \ldots \\ \vdots & \vdots & \ddots\end{pmatrix}\mathbf{V}^T. \tag{39}$$

It can be easily shown that the solution to (29) is given by

$$\widetilde{\mathbf{Y}} \leftarrow +S_{K\eta/\rho}\left(\left[\operatorname{vech}\left(\widetilde{\mathbf{T}}_1\right) + \mathbf{U}_1, \ldots, \operatorname{vech}\left(\widetilde{\mathbf{T}}_K\right) + \mathbf{U}_K\right]\right).$$

A pseudo-code of the proposed JICE algorithm is given in the table Algorithm 1, where we denote

$$\mathbf{U} = [\mathbf{U}_1, \ldots, \mathbf{U}_K]. \tag{40}$$

Note that the main loop of the algorithm involves two stages in each iteration, corresponding to the estimation and structure learning steps, described above. On the first approximation stage the values of $\widetilde{\mathbf{T}}_k$-s are updated, which can be done in parallel and is denoted in the code by the **parfor** loop. On the second stage, the code gathers the new $\widetilde{\mathbf{T}}_k$ values and updates $\widetilde{\mathbf{Y}}$ and $\mathbf{U}$, which can be viewed as the structure learning step.

Algorithm 1: The ADMM Implementation of the JICE

Input: $\mathbf{S}_k,\ \rho,\ \eta,\ \varepsilon,$
Output: $\widehat{\mathbf{T}}_k,$
Initialization : $\mathbf{U} \leftarrow 0,\ f_{\text{previous}} = +\infty,\ f_{\text{current}} = 0$
repeat
  parfor $k = 1\colon K$ **do**
    decompose:
    $\rho \operatorname{mat}\left(\widehat{\mathbf{Y}}_{:,k} - \mathbf{U}_k\right) - \mathbf{S}_k = \mathbf{D}_k\mathbf{\Lambda}_k\mathbf{D}_k^T,$
    $\widehat{\mathbf{T}}_k \leftarrow \frac{1}{2\rho}\mathbf{D}_k^T\left(\mathbf{\Lambda}_k + \sqrt{\mathbf{\Lambda}_k^2 + 4\rho\mathbf{I}}\right)\mathbf{D}_k,$
  end parfor
  $\widehat{\mathbf{Y}} \leftarrow$
  $S_{K\eta/\rho}\left(\left[\operatorname{vech}\left(\widehat{\mathbf{T}}_1\right) + \mathbf{U}_1, .., \operatorname{vech}\left(\widehat{\mathbf{T}}_K\right) + \mathbf{U}_K\right]\right),$
  $\mathbf{U} \leftarrow \mathbf{U} + \left[\operatorname{vech}\left(\widehat{\mathbf{T}}_1\right), \ldots, \operatorname{vech}\left(\widehat{\mathbf{T}}_K\right)\right] - \widehat{\mathbf{Y}},$
  $f_{\text{previous}} \leftarrow f_{\text{current}},$
  $f_{\text{current}} \leftarrow f(\widehat{\mathbf{T}}_1, \ldots, \widehat{\mathbf{T}}_K, \widehat{\mathbf{Y}}),$
until $|f_{\text{current}} - f_{\text{previous}}| \geqslant \varepsilon$

Bias Removal Step
decompose: $\widehat{\mathbf{Y}} = \widehat{\mathbf{U}}\widehat{\mathbf{Z}},$
parfor $k = 1\colon K$ **do**
  $\widehat{\mathbf{T}}_k \leftarrow \operatorname{argmin}_{\widetilde{\mathbf{T}}_k \subset \operatorname{span}(\widehat{\mathbf{U}})} -\log|\widetilde{\mathbf{T}}_k| + \operatorname{Tr}\left(\mathbf{S}_k\widetilde{\mathbf{T}}_k\right).$
end parfor

Under reasonable conditions, the general ADMM algorithm is claimed [28] to obtain the following convergence properties: the dual variable and the objective converge to their true values, whether convergence of the primal variables is not guaranteed. Under additional assumptions, primal variables converge to the solution, as well. In addition, as it is mentioned in [28], the ADMM can be quite slow to converge to high accuracy. Fortunately, in our problem modest accuracy is sufficient to obtain a good estimator. Our numerical simulations demonstrated that the most reliable stopping criterion is based on the convergence of the objective function value, which is implemented in Algorithm 1.

Another important question is the choice of the penalty parameter ρ. In our implementation we always took it constant $\rho = 1$, but many different approaches of tuning this parameter exist in literature. See [28] for a more full treatment of this issue and a list of references.

E. Upper Performance Bound

In this section we provide a high probability performance bound for the proposed JICE algorithm. In the derivation presented here we follow the technique used by [1], [29]–[31]. The method consists in showing that the extremal point of the target (22) lies within some small vicinity of the true parameter value with high probability. At first, the target is approximated by Taylor's polynomial with a second order remainder in the Lagrange form. Then concentration of measure phenomenon

is utilized to obtain high probability bounds on the derivatives and tune the regularization parameter. Finally, it is shown that the regularized solution must belong to some small ball around the true parameter value. The radius of this ball plays the role of the desired error bound with the probability guaranteed by the first step. Denote

$$\mathbf{Q} = \left[\mathrm{vec}\left(\mathbf{Q}_1\right), \ldots, \mathrm{vec}\left(\mathbf{Q}_K\right)\right], \tag{41}$$

and

$$\mathbf{S} = \left[\mathrm{vec}\left(\mathbf{S}_1\right), \ldots, \mathrm{vec}\left(\mathbf{S}_K\right)\right], \tag{42}$$

where $\mathbf{S}_k,\ k = 1, \ldots, K$ are the SCM-s defined in (21).

Theorem 1. *Assume that $n \geqslant p$ and $\eta \geqslant \frac{\|\mathbf{S}-\mathbf{Q}\|_2}{K}$, then with probability at least $1 - K\left(2^{-n+p-1} - e^{-c_1 n}\right)$,*

$$\left\|\widehat{\mathbf{Y}} - \mathbf{Y}\right\|_F \leqslant \frac{4\overline{\lambda}^2\sqrt{2r}K\eta}{c_2(1-\sqrt{(p-1)/n})^4}, \tag{43}$$

where c_1 and c_2 do not depend on K, p, n or r.

Proof. The proof can be found in the Appendix. □

As expected, this bound shows that the Frobenius norm of the error is proportional to the squared root of the rank rather than the dimension.

F. Regularization Parameter Choice

One of the main ingredients of the algorithm and the bound guaranteed by Theorem 1 is the condition on the regularization parameter η. Its optimal choice yielding the best performance characteristics of the algorithm depends on unobservable quantities. Therefore, in general one might need to tune the regularization parameter through cross validation [32]. Below we show how it can be calculated based on the given measurements. Let us start by estimating the norm $\|\mathbf{S} - \mathbf{Q}\|_2$.

Lemma 2. *(Lemma 4 from [18]) With probability at least $1 - 2e^{-c_3 p}$,*

$$\|\mathbf{S} - \mathbf{Q}\|_2 \leqslant \frac{1}{\underline{\lambda}\sqrt{n}}\left(\sqrt{2K} + c_4\sqrt{p^2}\right), \tag{44}$$

where c_3 and c_4 do not depend on K, p, n or r.

Recall that $p^2 = 2l + o(l),\ l \to \infty$ to get an immediate corollary of Theorem 1.

Corollary 1. *If $n \geqslant p$ and $\eta \asymp \frac{\|\mathbf{S}-\mathbf{Q}\|_2}{K}$, with high probability*

$$\left\|\widehat{\mathbf{Y}} - \mathbf{Y}\right\|_F \leqslant \frac{4\overline{\lambda}^2\sqrt{2r}}{\underline{\lambda}\sqrt{n}(1-\sqrt{(p-1)/n})^4}\left(\sqrt{2K} + C\sqrt{l}\right). \tag{45}$$

We believe that in all the obtained results the requirement $n \geqslant p$ is too strong and can be relaxed. We are currently pursuing a way to make this intuition a precise proof.

Apparently, in the real world environment the matrix $\mathbf{Q}$ is not known and we need to compute the value of η from the available data. Note that

$$\underline{\lambda} = \min_k \lambda_p(\mathbf{T}_k) = \min_k \frac{1}{\|\mathbf{Q}_k\|_2} = \frac{1}{\max_k \|\mathbf{Q}_k\|_2}, \tag{46}$$

and can naturally be estimated as

$$\widehat{\underline{\lambda}} = \frac{\left(1+\sqrt{\frac{p}{n}}\right)^2}{\max_k \|\mathbf{S}_k\|_2}, \tag{47}$$

where the last relation comes from Theorem II.13 from [33]. Now Lemma 2 and Corollary 1 together with (47) suggest a natural choice of the regularization parameter value

$$\eta = \frac{\max_k \|\mathbf{S}_k\|_2}{\sqrt{n}K(1+\sqrt{\frac{p}{n}})^2}\left(\sqrt{2K} + \sqrt{l}\right). \tag{48}$$

Below we numerically demonstrate the performance characteristics of the JICE algorithm with this choice of η.

V. Numerical Simulations

A. Circulant Model

For our first numerical experiment we took the circulant model. An important property of circulant matrices is that they are diagonalized by a Digital Fourier Transform (DFT) basis and, therefore, the structure is preserved under the inversion which we utilize below. In order for a real covariance matrix $\mathbf{Q}$ to be circulant, it should have an odd dimension p and its spectrum must be of the form

$$\mathrm{spec}(\mathbf{Q}) = [\xi_{(p+1)/2}, \xi_1, \ldots, \xi_{(p-1)/2}, \xi_{(p-1)/2}, \ldots, \xi_1], \tag{49}$$

with arbitrary positive ξ_i-s and the order corresponding to the standard ordering of the DFT basis vectors (so that the same eigenvalues correspond to the conjugated eigenvectors). In our experiment we have generated $K = 50$ random i.i.d. positive definite circulant matrices as

$$\mathbf{Q}_k = 2\mathbf{I} + \mathbf{DFT}\ \mathrm{spec}(\mathbf{Q}_k)\ \mathbf{DFT}^H, \tag{50}$$

where the spectrum $\mathrm{spec}(\mathbf{Q}_k)$ was generated according to (49) with ξ_i i.i.d uniform over the interval $[-1, 1]$ and **DFT** is the unitary DFT basis matrix of dimension p. Figure 1a shows the empirical, averaged over 1000 experiments, **MSE**-s of a few estimators as functions of the number of samples n in groups. The TSVD curve demonstrates the performance of the Truncated SVD after inversion (or pseudoinversion) of the group sample covariances, as proposed in the beginning of section IV and follows the development in [18] applied to inverse covariances. The truncation threshold for the SVD was chosen according to the rule described in [18] in such a way, that the truncated signal would convey 90% of the original power. Our JICE algorithm and its bias removing modifications are named JICE and JICE_BR, respectively. For comparison we also plot the **MSE**-s of the inverse SCM (ISCM), its projection onto the known structure (Projection) and the **CRB** bound calculated in (18). The graph demonstrates that when the number of samples n in each of the groups is relatively small, both JICE and JICE_BR significantly outperform the competitors. In addition, as expected, the bias removal step improves the performance significantly and is quite close to the **CRB**.

The main goal of our next simulation study was to examine how good is the choice of the regularization parameter given by (48). For this purpose in Figure 1b we compare the performances of the JICE_BR algorithm in the same circulant setting

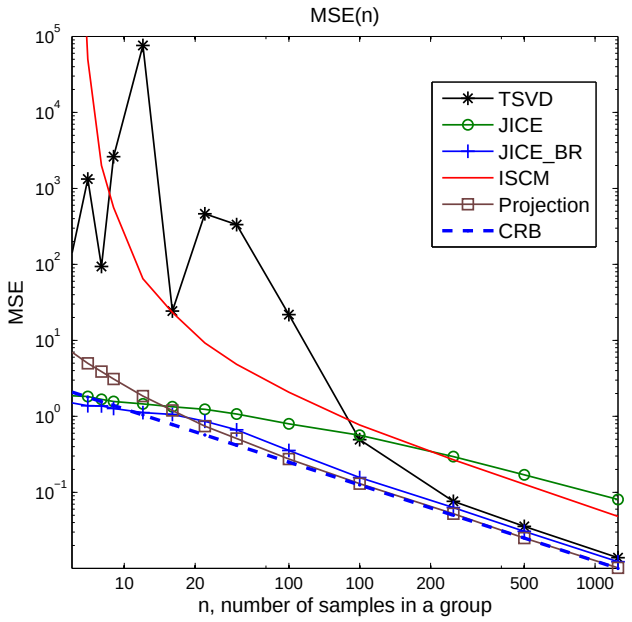

(a)

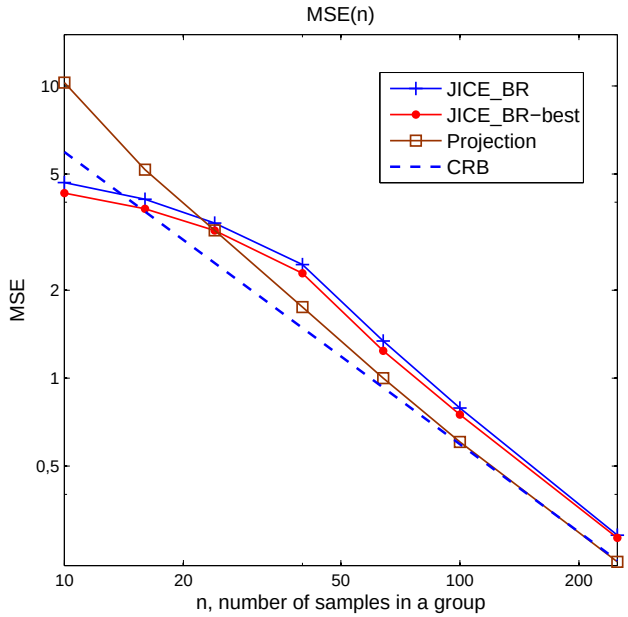

(b)

Fig. 1: JICE algorithm performance in the circulant structure case, $p = 5,\ l = 15,\ r = 5,\ K = 32$.

with the benchmark (JICE_BR-best). To get the JICE_BR-best curve we took a dense grid of regularization parameters $\mathcal{Z} = \frac{1}{\sqrt{nK}}[0.05,\ 0.1,\ 0.2,\ \ldots, 10,\ 25,\ 50]$ and, given the true covariances, chose the best $\eta_b \in \mathcal{Z}$ producing the smallest $\mathbf{MSE}$. These $\mathbf{MSE}$ as functions of n are depicted in the figure. As the Figure clearly suggests, the theoretically derived in (48) value of the regularization parameter demonstrates very good performance and can be readily used in applications.

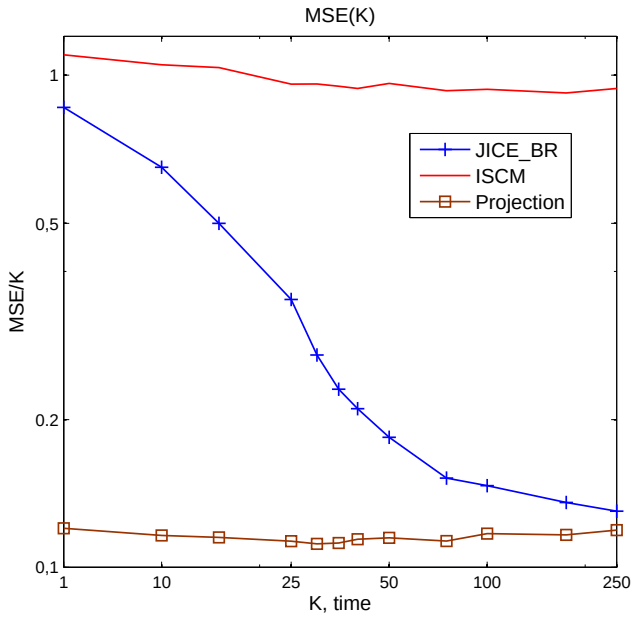

Fig. 2: Time-tracking learning process in Graphical Models, $p = 5,\ b = 1,\ l = 15,\ r = 9,\ n = 40$.

B. Graphical Models

The goal of this section is to examine the learning abilities of the JICE technique in Gaussian Graphical Models. The setting of the simulation identifies every other group of measurements $k = 1, \ldots, K$ with data coming with the successive clock ticks. The inverse covariances of the groups were generated by the random process

$$\bar{\mathbf{T}}_t = \mathbf{T}_{t-1} + \alpha \mathbf{\Delta}_t, \quad \mathbf{T}_t = \frac{\bar{\mathbf{T}}_t}{\left\|\bar{\mathbf{T}}_t\right\|}, \quad t = 0, 1, \ldots, K, \quad (51)$$

where α is a small weight and $\mathbf{\Delta}_t$ is such a random matrix preserving the sparsity pattern that $\mathbf{T}_t \succ \epsilon \mathbf{I},\ t = 0, 1, \ldots, K$. Specifically, we fixed $p = 5$ and chose the inverse covariances to be symmetric positive definite tridiagonal matrices with the smallest eigenvalue not less than $\varepsilon = 0.1$. The data generating algorithm worked as follows: on the fist step we generated symmetric tridiagonal $\mathbf{T}_0$ with uniformly over $[-1, 1]$ distributed entries and checked if $\lambda_p(\mathbf{T}_0) \geqslant \varepsilon$, if this condition was violated, we repeated this step until it was met. Then on each time count t we performed a similar procedure: we took symmetric tridiagonal $\mathbf{\Delta}_t$ with uniformly randomly over $[-1, 1]$ entries and checked if $\lambda_p(\mathbf{T}_t) \geqslant \varepsilon$. If the last condition was not fulfilled iteration was repeated until it was satisfied. In our experiment α was taken to be 0.25. The number of degrees of freedom in this model is $r = 9$ and the ambient dimension is $l = 15$. The number of measurements in each group was fixed to be $n = 40$. Figure 2 shows that the JICE_BR learns the structure with the time and effectively applies this knowledge in the estimation. All the results were averaged over 100 iterations.

VI. Conclusion

In this paper we consider the problem of joint inverse covariance estimation with linear structure, given heteroge-

neous measurements. The main challenge in this scenario is twofold. At first, the underlying structure is to be discovered and then it should be utilized to improve the concentrations estimation. We propose a novel algorithm coupling these two stages into one optimization program and propose its efficient implementation. The main aim of our current research is to provide tight upper bound guarantees on the performance of the proposed technique.

APPENDIX

Proof of Theorem 1. To simplify the proof, without loss of generality, we assume that $\mathbf{Y}$ and all its analogs are formed by $\mathrm{vec}\,(\mathbf{T}_k)$, rather than $\mathrm{vech}\,(\mathbf{T}_k)$.

Consider the first order Taylor's expansion of the target (22) with the remainder in Lagrange form in the vicinity of $\mathbf{Y}$

$$\begin{aligned}\Delta f &= f(\mathbf{Y}+\Delta\mathbf{Y}) - f(\mathbf{Y}) \\ &= \nabla_{\mathbf{Y}} f(\Delta\mathbf{Y}) + \frac{1}{2}\nabla^2_{\mathbf{Y}} f(\Delta\mathbf{Y}) + \eta\Delta\left\|\mathbf{Y}\right\|_*, \end{aligned} \tag{52}$$

where the derivatives are treated as linear and quadratic forms correspondingly, their subscripts denote the point of evaluation, and

$$\Delta\mathbf{Y} = [\mathrm{vec}\,(\Delta\mathbf{T}_1), \ldots, \mathrm{vec}\,(\Delta\mathbf{T}_K)], \tag{53}$$

$$\bar{\mathbf{Y}} = \mathbf{Y} + \alpha\Delta\mathbf{Y}, \;\; \alpha\in[0,1], \tag{54}$$

$$\Delta\left\|\mathbf{Y}\right\|_* = \left\|\mathbf{Y}+\Delta\mathbf{Y}\right\|_* - \left\|\mathbf{Y}\right\|_*. \tag{55}$$

Compute the derivatives

$$\begin{aligned}\nabla_{\mathbf{Y}} f(\Delta\mathbf{Y}) &= \frac{1}{K}\sum_{k=1}^{K}\mathrm{Tr}\left((\mathbf{S}_k - \mathbf{T}_k^{-1})\Delta\mathbf{T}_k\right) \\ &= \frac{1}{K}\mathrm{Tr}\left((\mathbf{S}-\mathbf{Q})^T\Delta\mathbf{Y}\right), \end{aligned} \tag{56}$$

$$\nabla^2_{\bar{\mathbf{Y}}} f(\Delta\mathbf{Y}) = \frac{1}{K}\sum_{k=1}^{K}\mathrm{Tr}\left(\left[(\mathbf{T}_k+\alpha\Delta\mathbf{T}_k)^{-1}\Delta\mathbf{T}_k\right]^2\right). \tag{57}$$

Consider the target function over the set satisfying

$$\lambda_p\left((\mathbf{T}_k+\alpha\Delta\mathbf{T}_k)^{-1}\right) \geqslant c_5\lambda_p(\mathbf{S}_k),\; k=1,\ldots,K. \tag{58}$$

Theorem 1.1 from [34] claims that when $n \geqslant p$, with probability at least $1-2^{-n+p-1}-e^{-c_1 n}$,

$$\begin{aligned}\lambda_p(\mathbf{S}_k) &\geqslant c_6\lambda_p(\mathbf{Q}_k)\frac{(\sqrt{n}-\sqrt{p-1})^2}{n} \\ &\geqslant \frac{c_6}{\overline{\lambda}}\left(1-\sqrt{\frac{p-1}{n}}\right)^2, \end{aligned} \tag{59}$$

where c_1 and c_6 are universal constants. The union bound implies that with probability at least $1-K\left(2^{-n+p-1}-e^{-c_1 n}\right)$ (59) holds simultaneously for all k and, therefore, after denoting $c_2 = c_5 c_6$, with the same probability

$$\lambda_p\left((\mathbf{T}_k+\alpha\Delta\mathbf{T}_k)^{-1}\right) \geqslant \frac{c_2}{\overline{\lambda}}\left(1-\sqrt{\frac{p-1}{n}}\right)^2. \tag{60}$$

over the set under investigation. Denote

$$\zeta = \frac{c_2}{\overline{\lambda}}\left(1-\sqrt{\frac{p-1}{n}}\right)^2, \tag{61}$$

and plug (60) into (57) to get

$$\nabla^2_{\bar{\mathbf{Y}}} f(\Delta\mathbf{Y}) \geqslant \frac{\zeta^2}{K}\sum_{k=1}^{K}\left\|\Delta\mathbf{T}_k\right\|_F^2 = \frac{\zeta^2}{K}\left\|\Delta\mathbf{Y}\right\|_F^2. \tag{62}$$

At the optimal point $\widehat{\mathbf{Y}}$, necessarily $\Delta f \leqslant 0$, thus

$$\frac{1}{K}\mathrm{Tr}\left((\mathbf{S}-\mathbf{Q})^T\Delta\mathbf{Y}\right) + \frac{\zeta^2}{2K}\left\|\Delta\mathbf{Y}\right\|_F^2 + \eta\Delta\left\|\mathbf{Y}\right\|_* \leqslant 0, \tag{63}$$

or

$$\frac{\zeta^2}{2}\left\|\Delta\mathbf{Y}\right\|_F^2 \leqslant -K\eta\Delta\left\|\mathbf{Y}\right\|_* - \mathrm{Tr}\left((\mathbf{S}-\mathbf{Q})^T\Delta\mathbf{Y}\right). \tag{64}$$

Introduce the SVD of $\mathbf{Y}$

$$\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = [\mathbf{U}_r\mathbf{U}_r^{\perp}]\mathbf{\Sigma}[\mathbf{V}_r\mathbf{V}_r^{\perp}]^T, \tag{65}$$

where $\mathbf{U}\in\mathbb{R}^{p^2\times p^2}, \mathbf{V}\in\mathbb{R}^{K\times K}$ are orthogonal, $\mathbf{U}_r\in\mathbb{R}^{p^2\times r}, \mathbf{V}_r\in\mathbb{R}^{K\times r}$ and the upper left part of the diagonal of $\mathbf{\Sigma}\in\mathbb{R}^{p^2\times K}$ contains the r non zero singular values of $\mathbf{Y}$. Let

$$\mathbf{\Gamma} = \mathbf{U}^T\Delta\mathbf{Y}\mathbf{V} = \begin{pmatrix}\mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12}\\ \mathbf{\Gamma}_{21} & \mathbf{\Gamma}_{22}\end{pmatrix}, \tag{66}$$

where $\mathbf{\Gamma}_{11}\in\mathbb{R}^{r\times r}$ and define

$$\Delta\mathbf{Y}'' = \mathbf{U}\begin{pmatrix}0 & 0\\ 0 & \mathbf{\Gamma}_{22}\end{pmatrix}\mathbf{V}^T, \tag{67}$$

$$\Delta\mathbf{Y}' = \Delta\mathbf{Y} - \Delta\mathbf{Y}''. \tag{68}$$

Lemma 3. *With the notation (67), (68),*

$$-\Delta\left\|\mathbf{Y}\right\|_* \leqslant \left\|\Delta\mathbf{Y}'\right\|_* - \left\|\Delta\mathbf{Y}''\right\|_*. \tag{69}$$

Proof. The proof can be found below. □

Plug (69) into (64) and use Holder's inequality

$$\mathrm{Tr}\left((\mathbf{S}-\mathbf{Q})^T\Delta\mathbf{Y}\right) \leqslant \left\|\mathbf{S}-\mathbf{Q}\right\|_2\left\|\Delta\mathbf{Y}\right\|_*, \tag{70}$$

to obtain

$$\begin{aligned}&\frac{\zeta^2}{2}\left\|\Delta\mathbf{Y}\right\|_F^2 \leqslant -K\eta\Delta\left\|\mathbf{Y}\right\|_* - \mathrm{Tr}\left((\mathbf{S}-\mathbf{Q})^T\Delta\mathbf{Y}\right) \\ &\leqslant K\eta(\left\|\Delta\mathbf{Y}'\right\|_* - \left\|\Delta\mathbf{Y}''\right\|_*) + \left\|\mathbf{S}-\mathbf{Q}\right\|_2(\left\|\Delta\mathbf{Y}'\right\|_* + \left\|\Delta\mathbf{Y}''\right\|_*), \end{aligned} \tag{71}$$

where in the last transition we have applied the triangle inequality. Recall that $\eta \geqslant \left\|\mathbf{S}-\mathbf{Q}\right\|_2/K$ to obtain

$$\frac{\zeta^2}{2}\left\|\Delta\mathbf{Y}\right\|_F^2 \leqslant 2K\eta\left\|\Delta\mathbf{Y}'\right\|_*. \tag{72}$$

Note that

$$\begin{aligned}\mathrm{rank}(\Delta\mathbf{Y}') &= \mathrm{rank}\begin{pmatrix}\mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12}\\ \mathbf{\Gamma}_{21} & 0\end{pmatrix} \\ &\leqslant \mathrm{rank}\begin{pmatrix}\mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12}\\ 0 & 0\end{pmatrix} + \mathrm{rank}\begin{pmatrix}\mathbf{\Gamma}_{11} & 0\\ \mathbf{\Gamma}_{21} & 0\end{pmatrix} \leqslant 2r, \end{aligned} \tag{73}$$

and use the following norm relation

$$\left\|\Delta\mathbf{Y}'\right\|_* \leqslant \sqrt{\mathrm{rank}(\Delta\mathbf{Y}')}\left\|\Delta\mathbf{Y}'\right\|_F, \tag{74}$$

together with the decomposition (67)-(68) to get

$$\frac{\zeta^2}{2}\left\|\Delta\mathbf{Y}\right\|_F^2 \leqslant 2K\eta\left\|\Delta\mathbf{Y}'\right\|_* \leqslant 2K\eta\sqrt{2r}\left\|\Delta\mathbf{Y}'\right\|_F \leqslant 2\sqrt{2r}K\eta\left\|\Delta\mathbf{Y}\right\|_F. \quad (75)$$

We finally achieve the desired bound

$$\left\|\Delta\mathbf{Y}\right\|_F \leqslant \frac{4\sqrt{2r}K\eta}{\zeta^2} = \frac{4\overline{\lambda}^2\sqrt{2r}K\eta}{c_2^2(1-\sqrt{(p-1)/n})^4}. \quad (76)$$

□

Proof of Lemma 3.

$$\left\|\mathbf{Y}+\Delta\mathbf{Y}''\right\|_* = \left\|\mathbf{U}_r\mathbf{U}_r^T\mathbf{Y}\mathbf{V}_r\mathbf{V}_r^T+\Delta\mathbf{Y}''\right\|_* = \left\|\mathbf{Y}\right\|_*+\left\|\Delta\mathbf{Y}''\right\|_*. \quad (77)$$

Together with the triangle inequality this yield

$$\begin{aligned}\left\|\mathbf{Y}+\Delta\mathbf{Y}\right\|_* &= \left\|(\mathbf{Y}+\Delta\mathbf{Y}'')+\Delta\mathbf{Y}'\right\|_* \\ &\geqslant \left\|\mathbf{Y}+\Delta\mathbf{Y}''\right\|_*-\left\|\Delta\mathbf{Y}'\right\|_* \\ &\geqslant \left\|\mathbf{Y}\right\|_*+\left\|\Delta\mathbf{Y}''\right\|_*-\left\|\Delta\mathbf{Y}'\right\|_*, \end{aligned} \quad (78)$$

and, finally,

$$-\Delta\left\|\mathbf{Y}\right\|_* = -\left\|\mathbf{Y}+\Delta\mathbf{Y}\right\|_*+\left\|\mathbf{Y}\right\|_* \leqslant \left\|\Delta\mathbf{Y}'\right\|_*-\left\|\Delta\mathbf{Y}''\right\|_*. \quad (79)$$

□

REFERENCES

[1] P. J. Bickel and E. Levina, "Regularized estimation of large covariance matrices," *The Annals of Statistics*, pp. 199–227, 2008.

[2] A. Kavcic and J. M. F. Moura, "Matrices with banded inverses: Inversion algorithms and factorization of Gauss-Markov processes," *IEEE transactions on Information Theory*, vol. 46, no. 4, pp. 1495–1509, 2000.

[3] A. Asif and J. M. F. Moura, "Block matrices with L-block-banded inverse: inversion algorithms," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 630–642, 2005.

[4] A. Dembo, C. Mallows, and L. Shepp, "Embedding nonnegative definite Toeplitz matrices in nonnegative definite circulant matrices, with application to covariance estimation," *IEEE Transactions on Information Theory*, vol. 35, no. 6, pp. 1206–1212, 1989.

[5] T. T. Cai, Z. Ren, and H. H. Zhou, "Optimal rates of convergence for estimating Toeplitz covariance matrices," *Probability Theory and Related Fields*, vol. 156, no. 1-2, pp. 101–143, 2013.

[6] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.

[7] O. Banerjee, L. El Ghaoui, and A. d'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data," *The Journal of Machine Learning Research*, vol. 9, pp. 485–516, 2008.

[8] J. Guo, E. Levina, G. Michailidis, and J. Zhu, "Joint estimation of multiple graphical models," *Biometrika*, vol. 98, no. 1, pp. 1–15, 2011.

[9] O. Besson, S. Bidon, and J.-Y. Tourneret, "Covariance matrix estimation with heterogeneous samples," *Signal Processing, IEEE Transactions on*, vol. 56, no. 3, pp. 909–920, 2008.

[10] S. Bidon, O. Besson, and J.-Y. Tourneret, "A Bayesian approach to adaptive detection in nonhomogeneous environments," *Signal Processing, IEEE Transactions on*, vol. 56, no. 1, pp. 205–217, 2008.

[11] A. Aubry, V. Carotenuto, A. De Maio, and G. Foglia, "Exploiting multiple a priori spectral models for adaptive radar detection," *Radar, Sonar & Navigation, IET*, vol. 8, no. 7, pp. 695–707, 2014.

[12] W. Lee and Y. Liu, "Joint estimation of multiple precision matrices with common structures," *Journal of Machine Learning Research*, vol. 16, pp. 1035–1062, 2015.

[13] P. Danaher, P. Wang, and D. M. Witten, "The joint graphical lasso for inverse covariance estimation across multiple classes," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 76, no. 2, pp. 373–397, 2014.

[14] A. Wiesel, O. Bibi, and A. Globerson, "Time varying autoregressive moving average models for covariance estimation," *IEEE Transactions on Signal Processing*, vol. 61, no. 11, pp. 2791–2801, 2013.

[15] E. Moulines, P. Priouret, F. Roueff *et al.*, "On recursive estimation for time varying autoregressive processes," *The Annals of statistics*, vol. 33, no. 6, pp. 2610–2654, 2005.

[16] A. Ahmed and E. P. Xing, "Recovering time-varying networks of dependencies in social and biological studies," *Proceedings of the National Academy of Sciences*, vol. 106, no. 29, pp. 11 878–11 883, 2009.

[17] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[18] I. Soloveychik and A. Wiesel, "Joint covariance estimation with mutual linear structure," *arXiv preprint arXiv:1507.00123*, 2015.

[19] ——, "Group symmetric robust covariance estimation," *arXiv preprint arXiv:1412.2345*, 2014.

[20] S. L. Lauritzen, "Graphical models," *Oxford University Press*, 1996.

[21] A. Wiesel, Y. C. Eldar, and A. O. Hero, "Covariance estimation in decomposable Gaussian graphical models," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1482–1492, 2010.

[22] Y.-H. Li and P.-C. Yeh, "An interpretation of the Moore-Penrose generalized inverse of a singular Fisher Information Matrix," *IEEE Transactions on Signal Processing*, vol. 60, no. 10, pp. 5532–5536, 2012.

[23] G. Tang and A. Nehorai, "Lower bounds on the mean-squared error of low-rank matrix reconstruction," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4559–4571, 2011.

[24] E. J. Candes and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational mathematics*, vol. 9, no. 6, pp. 717–772, 2009.

[25] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.0 beta," http://cvxr.com/cvx, Sep. 2013.

[26] ——, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.

[27] I. Soloveychik and A. Wiesel, "Tyler's covariance matrix estimator in elliptical models with convex structure," *IEEE Transactions on Signal Processing*, vol. 62, no. 20, pp. 5251–5259, 2014.

[28] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[29] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu, "High-dimensional covariance estimation by minimizing l_1-penalized log-determinant divergence," *Electronic Journal of Statistics*, vol. 5, pp. 935–980, 2011.

[30] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu, "Sparse permutation invariant covariance estimation," *Electronic Journal of Statistics*, vol. 2, pp. 494–515, 2008.

[31] I. Soloveychik and A. Wiesel, "Performance analysis of Tyler's covariance estimator," *IEEE Transactions on Signal Processing*, vol. 63, no. 2, pp. 418–426, 2015.

[32] P. J. Bickel, B. Li, A. B. Tsybakov, S. A. van de Geer, B. Yu, T. Valdés, C. Rivero, J. Fan, and A. van der Vaart, "Regularization in statistics," *Test*, vol. 15, no. 2, pp. 271–344, 2006.

[33] K. R. Davidson and S. J. Szarek, "Local operator theory, random matrices and Banach spaces," *Handbook of the geometry of Banach spaces*, vol. 1, no. 317-366, p. 131, 2001.

[34] M. Rudelson and R. Vershynin, "Smallest singular value of a random rectangular matrix," *Communications on Pure and Applied Mathematics*, vol. 62, no. 12, pp. 1707–1739, 2009.